arXiv:1408.6288v1 [stat.CO] 27 Aug 2014

Decreasing flow uncertainty in Bayesian inverse problems through Lagrangian drifter control

D. McDougall[†] C. K. R. T. Jones[‡]

August 28, 2014

Abstract

Commonplace in oceanography is the collection of ocean drifter positions. Ocean drifters are devices that sit on the surface of the ocean and move with the flow, transmitting their position via GPS to stations on land. Using drifter data, it is possible to obtain a posterior on the underlying flow. This problem, however, is highly underdetermined. Through controlling an ocean drifter, we attempt to improve our knowledge of the underlying flow. We do this by instructing the drifter to explore parts of the flow currently uncharted, thereby obtaining fresh observations. The efficacy of a control is determined by its effect on the variance of the posterior distribution. A smaller variance is interpreted as a better understanding of the flow. We show a systematic reduction in variance can be achieved by utilising controls that allow the drifter to navigate new or interesting flow structures, a good example of which are eddies.

1 Introduction

The context of the problem we address in this paper is that of reconstructing a flow field from Lagrangian observations. This is an identical twin experiment in which a true flow field is unknown but from which Lagrangian type observations are extracted. It is assumed that little is known about the functional form of the flow field except that is barotropic, incompressible and either steady or with simple known time dependence. Note that, since it is incompressible and two-dimensional (barotropic), the field is given by a stream function $\psi(x, y, t)$. The objective is then to reconstruct an estimate of this stream function from Lagrangian observations, along with an associated uncertainty.

The question addressed is whether the uncertainty of the reconstruction can be reduced by strategic observations using Lagrangian type instruments. The measuring devices are assumed to be controllable and their position can

[†]Institute for Computational Engineering and Sciences, University of Texas at Austin, TX, USA

[‡]Mathematics Department, University of North Carolina at Chapel Hill, NC, USA

be registered at appropriate time intervals. Since the control is known, being prescribed by the operator, it is reasonable to believe that information can be garnered from the position observations. The issue is whether we can improve the information content in these observations by controlling the instruments to move into specific flow regimes.

Estimating ocean flows has a long history. First, a comparison of model forecast errors in a barotropic open ocean model can be found in [49], with emphasis on how forecasts are sensitive to boundary information. Application of the Kalman filter with Lagrangian observations can be seen as early as 1982 [7, 41, 43, 11]. For a variational least-squares approach to eddy estimation, the reader is directed to [50]. A standard mathematical framework for incorporating Lagrangian observations appeared in 2003 [29]. Finally, [45] exposes a novel approach to ocean current observations involving the treatment of sea turtles as Lagrangian observers. A good overview of some operational ocean apparatus can be found in [51].

The underlying philosophy of our approach is that mesoscale ocean flow fields are dominated by coherent features, such as jets and eddies. If the instrument can be controlled to move into and through these structures then the information gained should be richer in terms of capturing the key properties of the flow field.

Of course, there is some circularity inherent in this approach; we want to get the key features of the flow field but need to know them in order to control the vehicle toward them. We first take a "proof of concept" approach and see if we follow a simple strategy that we happen to know takes into another eddy, as opposed to one that doesnt, then the uncertainty is reduced. We then postulate a good way of developing a control based purely on local information. The idea is to use the known local information of the flow field, from reconstructing the flow using observations up to a certain point in time, to form a control that takes the instrument away from the eddy it is currently stuck in.

To obtain the flow reconstruction, one needs to solve a Bayesian inverse problem [28]. There are numerous ways to solve Bayesian inverse probems, with the core methods being Kalman filtering and smoothing [26, 27, 52, 18, 22, 1]; variational methods [37, 8, 35, 36, 32, 54, 16, 30, 31]; particle filtering [17, 34]; and sampling methods [15, 12, 13, 33, 4, 2, 3, 20, 25, 42, 46, 47, 48, 9, 39, 19, 6, 5]. The resulting solution to a Bayesian inverse problem is a probability distribution, called the posterior distribution, over some quanitity of interest from which one can compute estimates with associated uncertainties. Bayesian inverse problems enable well-informed predictions.

The paper is organised in the following manner. The second section sets up the Bayesian inverse problem and specifies all the assumptions in the prior and likelihood distributions. The third section does applies two flow-independent (naïve) controls, a zonal (East-West) control, and a bidirectional (North-Easterly) control. These are applied to both the perturbed and unperturbed flows. We measure performance of the addition of each by looking at the posterior variance on the velocity field and show two main results. When the fluid flow drifter is trapped in a recirculation regime, the magnitude of the control is the main player in pushing the drifter out of the eddy. We show that, for the unperturbed

flow, when the control magnitude is large enough a significant reduction in the posterior variance is achieved. In the perturbed flow, we show robustness of the posterior variance with respect to the perturbation parameter. More specifically, its structure as a function of control magnitude is carried over from the time-independent flow model. Moreover, we observe an additional, and separate, decrease in posterior variance as a function of control magnitude corresponding to the purely time-dependent part of the flow. The fourth section examines the use of an a posteriori control, a control calculated using information from a previous Bayesian inversion done with no control present. Here the control magnitude corresponds geometrically to the distance between the drifter and a hyperbolic fixed point of an eddy transport barrier in the flow. As the control magnitude increases, the drifter gets closer to a hyperbolic fixed point of the drifter evolution equation and, for the unperturbed flow, a substantial decrease in posterior variance is observed. Hyperbolic fixed points of the drifter equations join transport barriers in the flow and act as a boundary to observations. Observing near these points outweighs the negative effects produced by polluting the observations with a large control size relative to the size of the flow. This gives a novel geometric correspondence between the control utilised here and the structure of the posterior variance as a function of control magnitude. The fifth section concludes the paper.

2 Setup

We begin by prescribing the stream function of the flow field the drifters will move in. We will call this flow field the 'truth' and later we try to reconstruct it from noisy observations. The truth flow we will use is an explicit solution to the barotropic vorticity equations [44],

$$\psi(x, y, t) = -cy + A\sin(2\pi kx)\sin(2\pi y) + \varepsilon\psi_1(x, y, t),$$

on the two dimensional torus $(x, y) \in \mathbb{T}^2$, where the perturbation we will use is given by

$$\psi_1(x, y, t) = \sin(2\pi x - \pi t)\sin(4\pi y).$$

The corresponding flow equation is as follows

$$\frac{\partial \mathbf{v}}{\partial t} = \varepsilon\partial_t\nabla^\perp\psi_1, \quad t > 0 \tag{1}$$

We will explore two cases. The first case is when $\varepsilon = 0$ and the underlying flow is steady. The second case is when $\varepsilon \neq 0$ and the time-dependent perturbation smears the underlying flow in the x-direction. Drifters placed in the flow $\mathbf{v}$ will obey

$$\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = \mathbf{v}(\mathbf{x}, t) + \mathbf{f}(\mathbf{x}, t).$$

The function $\mathbf{f}$ is called the *control*, the choice of which requires explicit diction. We consider two cases of control: a) flow-independent; and b) a posteriori. Flow-independent controls, are controls that do not systematically utilise information

regarding the underlying flow, $\mathbf{v}$. A posteriori controls harness information from a previous Bayesian update. Our soup-to-nuts methodology for assessing the efficacy for each case of control is as follows. First, drifter dynamics are obtained by solving

$$\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = \mathbf{v}(\mathbf{x}, t), \quad 0 < t \leq t_{K/2} \tag{2}$$

$$\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = \mathbf{v}(\mathbf{x}, t) + \mathbf{f}(\mathbf{x}), \quad t_{K/2} < t \leq t_K, \tag{3}$$

where $\mathbf{v}$ solves (1). Then observations of the drifter locations $\mathbf{x}$ are collected into an observation vector for both the controlled and uncontrolled parts

$$\begin{aligned} \mathbf{y}_k^1 &= \mathbf{x}(t_k) + \eta_k, \quad \eta_k \sim \mathcal{N}(0, \sigma^2 I_2), \quad k = 1, \ldots, K, \\ \rightsquigarrow \mathbf{y} &= \mathcal{G}(\mathbf{v}_0) + \eta, \quad \eta \sim \mathcal{N}(0, \sigma^2 I_{2K}), \end{aligned} \tag{4}$$

where $\mathbf{v}_0$ is the initial condition of the model (1). The map $\mathcal{G}$ is called the *forward operator* and maps the object we wish to infer to the space in which observations are taken.

Flow-independent controls $\mathbf{f}$ are independent of $\mathbf{y}^1$. We will utilise two such controls: a time-independent zonal control $\mathbf{f}(\zeta, 0)$; and a time-independent bi-directional control $\mathbf{f}(\zeta, \zeta)$. The a posteriori control we execute is one that forces drifter paths to be non-transverse to streamlines of the underlying flow. Namely, $\mathbf{f}(\mathbf{x}) = -\zeta \nabla \mathbb{P}(\psi_0 | \mathbf{y}^1)$, where $\psi_0(\mathbf{x}) = \psi(\mathbf{x}, 0)$. Our aim is to understand the effect of the control magnitude ζ and the resulting drifter path on the posterior distribution over the initial condition of the model $\mathbb{P}(\mathbf{v}_0 | \mathbf{y}^1, \mathbf{y}^2)$.

Encompassing our beliefs about how the initial condition, $\mathbf{v}_0$, should look into a prior probability measure, μ_0, it is possible to express the posterior distribution in terms of the prior and the data using Bayes's theorem. Bayes's theorem posed in an infinite dimensional space says that the posterior probability measure on $\mathbf{v}_0$, $\mu^{\mathbf{y}}$, is absolutely continuous with respect to the prior proability measure [53]. Furthermore, the Radon-Nikodym derivative between them is given by the likelihood distribution of the data,

$$\frac{\mathrm{d}\mu^{\mathbf{y}}}{\mathrm{d}\mu_0}(\mathbf{v}_0) = \frac{1}{Z(\mathbf{y})} \exp\left(\frac{1}{2\sigma^2} \|\mathcal{G}(\mathbf{v}_0) - \mathbf{y}\|^2\right),$$

where the operator $\mathcal{G}$ is exactly the forward operator as described in (4) and $Z(\mathbf{y})$ is a normalising constant. We utilise a Gaussian prior measure on the flow initial condition, $\mu_0 \sim \mathcal{N}(0, (-\Delta)^{-\alpha})$. For our purposes, we choose $\alpha = 3$ so that draws from the prior are almost surely in the Sobolev space $H^1(\mathbb{T}^2)$ [53, 10]. The posterior is a high dimensional non-Gaussian distribution requiring careful probing by use of a suitable numerical method. The reader is referred to [53] for a full and detailed treatment of Bayesian inverse problems on function spaces.

To solve the above Bayesian inverse problem, we use a Markov chain Monte Carlo (MCMC) method. MCMC methods are a class of computational techniques for drawing samples from a unknown target distribution. Throughout this paper, we have chosen to sample the posterior using a random walk

Metropolis-Hastings method on function space [14, 15, 10]. Using this approach, one can draw samples from the posterior distribution, obtaining its shape exactly. This is of use when the posterior distribution is not a Gaussian and cannot be uniquely determined by its first and second moments. The application of MCMC methods to solve Bayesian inverse problems is widespread. For examples of their use, see [12, 13, 33, 4, 2, 3, 21, 20, 38, 40, 25, 24, 42]. The theory above is all done in an infinite dimensional setting. Numerically and operationally, a finite dimensional approximation is made. In the case of the Karhunen-Loève expansion this approximation is done by truncation. A choice must be made in where to truncate, and this choice coincides with a modelling assumption—there are no frequencies of order larger than the truncation wavenumber. If it is feasible that solutions to the inverse problem do in fact admit higher-order frequencies, it is necessary to rethink this assumption. Throughout this paper the data and initial conditions are known and the truncation is chosen to be much larger than necessary to mitigate the effects of poor modelling assumptions. Practically, the true solution to one's problem is unknown. In this scenario, care and diligence are necessary in choosing appropriate prior assumptions.

3 Results: flow-independent control

Flow-independent controls concern the influence of an ocean drifter without using knowledge of the underlying flow. They are constructed in such a way as to be independent of $\mathbf{y}^1$.

3.1 Zonal control

Figure 1(a) shows the variance of the horizontal component of the flow as a function of control magnitude in the max norm, the L^1 norm and L^2 norm. The horizontal axis denotes the strength of the control. The vertical black dotted line corresponds to a critical value for the magnitude. Values of ζ less than this correspond to controls not strong enough to force the drifter out of the eddy. Conversely, values bigger correspond to controls that push the drifter out of the eddy.

Experiments were done for $\zeta = 0, 0.25, 0.5, \ldots, 2.75, 3$. The case $\zeta = 1.75$ was the first experiment in which we observed the drifter leaving the recirculation regime. The black line shows the maximum value of the variance over the domain $[0, 1] \times [0, 0.5]$. The magenta line and cyan line show the L^2 norm and L^1 norm, respectively. The minimum value of the variance is small enough to be difficult to see on the plot but remains consistently small, so it has been omitted for clarity reasons. There are some notable points to make here. Firstly, above the critical value (where the drifter leaves the eddy) we see that the size of the variance decreases in all of our chosen norms. We have learned more about the flow around the truth by forcing the drifter to cross a transport boundary and enter a new flow regime. Secondly, below the critical region (where the drifter

does not leave the eddy) we see an initial increase in the size of the variance. There are many factors at play here. We will try to shed some light on them.

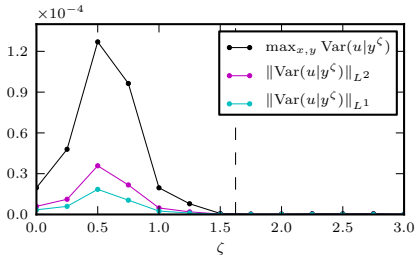

(a) The norm of the variance decreases as the glider is forced across the transport boundary and out of the eddy. The bump occurs as a result of the glider exploring a slow part of the flow.

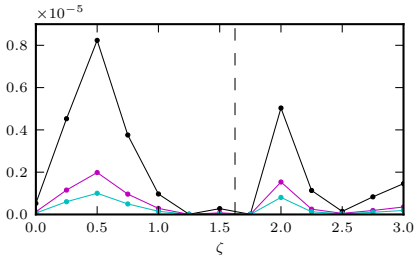

(b) The norm of the variance decreases as the glider is forced across the transport boundary and out of the eddy. The second bump appears because the glider re-enters a time-dependent eddy.

Figure 1: Posterior variance as a function of control magnitude, ζ, for (a) the time-independent model; and (b) the time-dependent model.

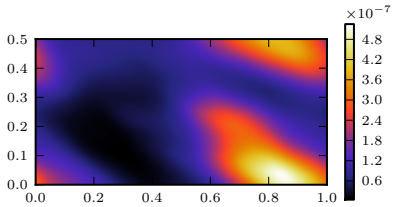

Figure 2: Horizontal component of the posterior variance for the case $\zeta = 0$. The black area in the lower left corresponds neatly with the region in which observations are taken.

For small ζ, the controlled and uncontrolled paths along which we take observations are close. Their closeness and the size of σ^2 creates a delicate interplay between whether they are statistically indistinguishable or not. If they are indistinguishable up to two or three standard deviations, this could explain the increase and then decrease of the variance below the critical value. Secondly, as ζ increases initially, the controlled path gets pushed down near the elliptic stagnation point of the flow (see figure 3). If this region is an area where the flow is smaller in magnitude than the flow along the uncontrolled path, this is equivalent to an increase in the magnitude of the control relative to the underlying flow. This leads to the observations becoming polluted by f.

Exploring this further, we compute the mean magnitude of the flow along the controlled path of the drifter. More formally, we solve (3) to obtain a set of

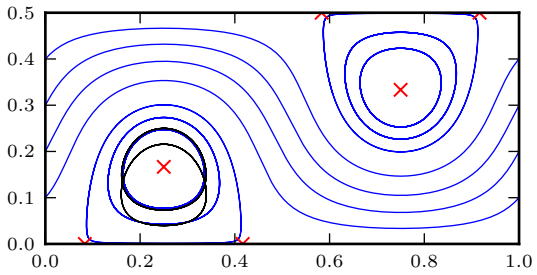

Figure 3: True glider path (black) for some positive ζ less than the critical value. Blue lines are streamlines of the true flow. Red crosses are zeros of the flow: fixed points of the passive glider equation.

points $\{zk = z(t_k)\}_{k=1}^{K}$. Then we compute the mean flow magnitude as follows

$$\langle v \rangle = \frac{2}{K} \sum_{k=K/2+1}^{K} v(z_k). \tag{5}$$

This quantity is computed for each fixed ζ the result is plotted in figure 4. The mean flow magnitude is given by the magenta line in this figure and the black dotted line depicts the flow magnitude. Notice the first three values of ζ which the mean flow magnitude decreases in. This is equivalent to and increase in the magnitude of the control relative to the magnitude of the underlying flow and so the information gain from taking observations here decreases. This corresponds nicely with the first three values of ζ in figure 1(a) that show an increase in variance. Notice also that for the other values of ζ the mean flow magnitude shows a mostly increasing trend, consistent with a decrease in the posterior variance.

Note that the region below the critical value correspond to control magnitudes that are too small to push the glider out of the eddy *in the unperturbed case* $\varepsilon = 0$. The region above the critical value corresponds to values of ζ for which the glider leaves the eddy, this is also in the unperturbed case. Experiments were done for $\zeta = 0, 0.25, 0.5, \ldots, 2.75, 3$. In the case $\varepsilon = 0$, the value $\zeta = 1.75$ was the first experiment in which we observed the glider leaving the recirculation regime. The black line shows the maximum value of the variance over the domain $[0, 1] \times [0, 0.5]$. The magenta line and cyan line show the L^2 norm and L^1 norm, respectively. There are some notable points to make. Firstly, below the critical magnitude (where the glider leaves the eddy in the unperturbed case) we see a sizeable reduction of posterior variance in the max norm as the critical magnitude is approached. To establish a connection in uncertainty quantification between the time independent and time-periodic case is of great scientific interest and that connection has been made evident here. Note that as ζ increases and progresses further into the region above the critical magnitude, the posterior variance repeats the increasing/decreasing structure induced by the eddy that we observed in the region below the critical control

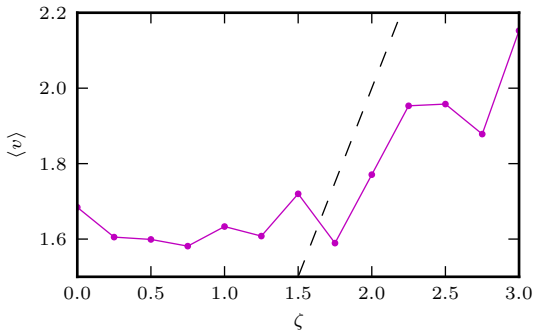

Figure 4: Mean magitude of the flow along the control path (purple) against the size of the control (black dashed line). When the gradient of the flow magnitude is large compared with that of the control magnitude, the posterior variance is small.

magnitude. The new effects introduced into this region are purely form the time-dependent nature of the moving eddy. The reason for their presence is much the same as in the time-independent case; observations trapped within an eddy regime.

We have learned more about the flow around the truth by forcing the glider into the meandering jet flow regime. The benefits of such a control occur at exactly the same place as in the time-independent case; as the drifter leaves the eddy in the unperturbed flow. However, extra care is required when the flow is time-dependent and the eddy moves. One cannot simply apply the same control techniques as is evidenced by the extra bump in variance in the region above the critical magnitude. Of particular use would be extra eddy-tracking information to construct an a posteriori control to keep the variance small.

3.2 Bi-directional control

Now we provide the analogue of figure 1(a) for the bi-directional forcing function. This is shown in figure 5(a). We see similar behaviour for the variance of the posterior distribution again. Below the critical magnitude, the values of ζ for which the drifter is not forced hard enough to leave the recirculation regime, we see an initial increase in the size of the posterior variance. Then we observe a decrease in posterior variance as ζ approaches a value large enough to push the drifter out of the eddy regime, the region above the critical value.

To explain the initial increase in the posterior variance below the critical magnitude, we calculate the mean flow magnitude just as in (5). This is shown in figure 6. We see an initial period where the mean flow along the controlled path remains almost constant. As a consequence, the magnitude of the forcing increases relative to the magnitude of the flow. This pollutes the observations and leads to an increased posterior variance just as we have observed in the

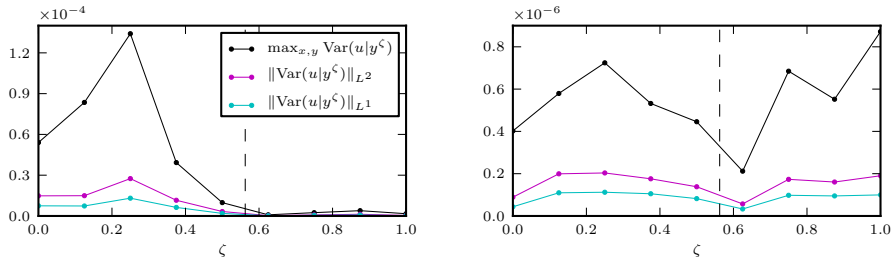

(a) The norm of the variance decreases as the glider is forced across the transport boundary and out of the eddy.

(b) The norm of the variance decreases as the glider is forced across the transport boundary and out of the eddy. The time-dependent part of the model pollutes the variance once the glider leaves the eddy.

Figure 5: Posterior variance as a function of control magnitude, ζ, for (a) the time-independent model; and (b) the time-dependent model.

previous section. We also see the opposite effect; the big jump in flow magnitude at $\zeta = 0.5$ (and consequently when the drifter escapes the gyre) is attested as the cause of the decrease in posterior variance as we enter the region above the critical flow magnitude of figure 5(a).

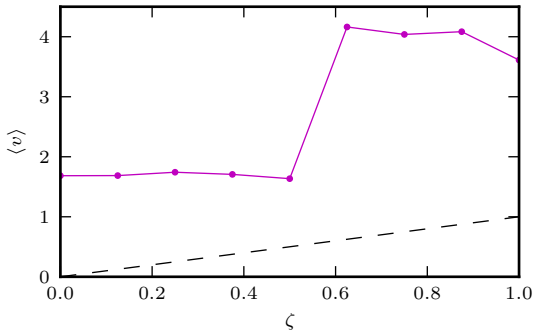

Figure 6: Mean magitude of the flow along the control path (purple) against the size of the control (black dashed line). When the gradient of the flow magnitude is large compared with that of the control magnitude, the posterior variance is small.

The cases of forcing explored thus far are $f(z) = (\zeta, 0)^\top$ and $f(z) = (\zeta, \zeta)^\top$. The main results are summarised by referring to figure 1(a) and figure 5(a). In these two cases, we see strikingly similar structure of the posterior variance as a function of control magnitude. The initial increase in posterior variance within the eddy; decreasing posterior variance as the flow path of the drifter approaches the transport boundary and small posterior variance (compared to the case $\zeta = 0$) once a new flow regime is being observed. Compare the values

of ζ for with this behaviour occurs. Notice that the values of ζ in figure 1(a) are about three times larger than those in figure 5(a). One factor at play here is the relative magnitude of the controls in each case. For $\zeta = 1$, the control has magnitude 1 in the zonal case, and magnitude $\sqrt{2}$ in the bi-directional case. Even scaling the results in the bi-directional case by $\sqrt{2}$, notice that the value of ζ for which the drifter first leaves the eddy, is $\zeta = \frac{\sqrt{2}}{2}$ and this is still smaller than $\zeta = 1.5$ for the x-directional case. The final factor affecting the scaling is the dynamics of the system after the forcing has been applied. Controlling in only the horizontal direction will require a larger magnitude force to push the drifter out of the eddy than when forcing in both the x and y directions simultaneously.

An analogue for figure 1(b) for the new forcing function is shown in figure 5(b). We see similar behaviour for the variance of the posterior distribution. Again, the region below the critical magnitude corresponds to values of ζ that are not big enough to push the drifter out of the recirculation regime in the *unperturbed* case. Just as in figure 1(b), we see the unperturbed eddy affecting the variance of the posterior distribution on the flow in the classic 'bump' fashion. We observe a reduction in posterior variance as ζ approaches a value large enough to push the glider out of the eddy regime (in the case $\varepsilon = 0$). In the region below the critical magnitude, the time-dependent flow effects take over and push the variance up. Again, a connection of uncertainty quantification is made between the time-independent case and the case where the flow is perturbed by a time-periodic disturbance, this connection lies entirely within the region below the critical control magnitude.

4 Results: a posteriori control

In section 3 we concluded that crossing a transport boundary and entering a new flow regime has the desirable effect of reducing the posterior variance. Crossing into new flow regimes with a stationary flow can be translated to travelling transversely against the streamlines of the underlying flow. For the recirculation regime located in the bottom-left area, particles in the fluid will move in an clockwise fashion. The gradient of the stream function will therefore point in towards the fixed point at $z = (1/4, 1/6)$. The negative gradient of the stream function points towards the fixed point at $z = (3/4, 1/3)$. Therefore, to escape the recirculation regime we choose,

$$f(z) = -\zeta \nabla_z (\mathbb{E}(\psi | y^1)), \tag{6}$$

for the controlled drifter model, where ψ is the stream function of the flow v. The rationale behind this choice is that, if the posterior mean stream function is a good estimator of the flow, the drifter will be forced transversely with the stream lines and escape the recirculation regime and allow us to make observations in a new flow regime.

Figure 7(a) depicts the variance of the horizontal component as the strength of the control, ζ, is varied. Note that we do not see the same behaviour as

we do for the two naïve controls chosen in section 3. We see a large band of values of ζ for which the posterior variance oscillates, leading to a lack of information gain in the knowledge of the flow. From about $\zeta = 0.5$ to $\zeta = 0.55$, we see a structurally significant reduction in posterior variance where we have a sustained gain in information about the underlying flow field. This is attributed to a drifter path that explores an 'interesting' part of the flow where a lot of information can be obtained from observations. To explore the geometric correspondence between the variance reduction for $\zeta = 0.5$ to $\zeta = 0.55$, we show figure 4. This figure presents the true path of the drifter for $\zeta = 0.3, \ldots, 0.55$. The light pink path corresponds to a value of $\zeta = 0.3$ and the purple path corresponds to $\zeta = 0.55$. Notice that as ζ increases, the true path forms a kink and forms a trajectory close to the zero of the flow at $(x, y) = (7/12, 1/2)$. Just as we have seen in section 3, we observe a transient period in the posterior variance until we utilise a control for which the true path explores new aspects of the flow compared with other 'nearby' controls. Interestingly, also note that we observe this reduction in variance despite the true path navigating near a zero of the flow, where we also satisfy the fact the the size of the control is large in comparison to the flow. In this case, a logical conclusion here would be that the information gain from observing near an interesting flow structure heavily outweighs the information loss in polluting the observations with such a control. The cost of polluting the observed data can be seen by computing the most structurally significant reduction in the posterior variance and comparing this with figure 5(a), for example. By 'most structurally significant' we loosely mean the most dramatic reduction that leads to the most benefit in knowledge of the underlying flow. In this example, this occurs between $\zeta = 0.52$ and $\zeta = 0.55$, where it is approximately 3×10^{-5}. In the case of the bi-directional control, where the relative size of the flow *increases* for the values of ζ that give a reduction in variance , it occurs between $\zeta = 0.25$ and $\zeta = 0.625$ where it is approximately 1.5×10^{-4}. This is about an order of magnitude bigger, crystallising the tradeoff between polluting the observed data versus exploring 'interesting' parts of the flow. If the posterior mean is a good estimator of the underlying flow, utilising a control of this nature is beneficial if the drifters navigates close to a hyperbolic fixed point of the passive drifter model equation.

The first thing to note is that we do not see the same behaviour as we do for the two naïve controls chosen in section 3. Nor do we see similar structures when compared with figure 7(a). For each value of ζ, it is the case that the true path navigates to the time-dependent eddy surrounding the zero of the flow at the point $(x, y) = (3/4, 1/3)$. The second thing to note is that for all of these values of control magnitude, the smaller values tend to do better than the larger ones The variance is lower in the cases $\zeta = 0.21$ and $\zeta = 0.27$ because the true path is navigating towards one of the hyperbolic fixed points of the eddy. A novel connection is established between the behaviour of these two controls in both the time-independent case and the time-periodic case.

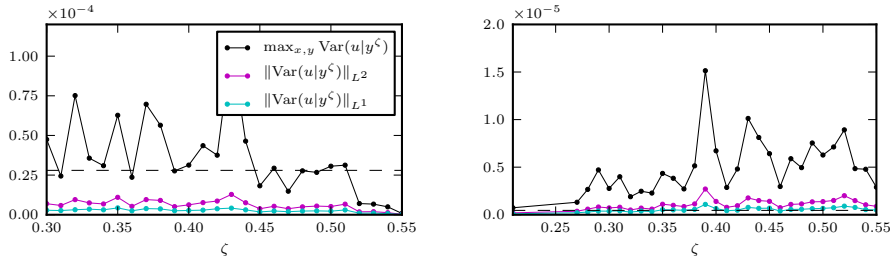

(a) The norm of the variance decreases as the glider is forced towards the a saddle point in the flow. No clear gain is made otherwise.

(b) No clear gain is made in the case of the time-dependent model.

Figure 7: Posterior variance as a function of control magnitude, ζ, for the a posteriori control in the case of: (a) a time-independent model; and (b) a time-dependent model.

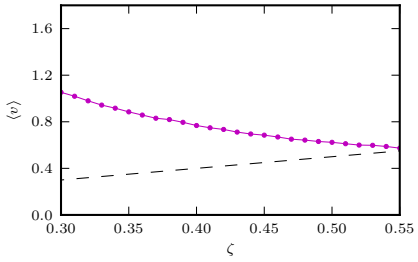

Figure 8: Mean magitude of the flow along the control path (purple) against the size of the control (black dashed line). Though the gradient of the flow magnitude is small compared with that of the control magnitude, the posterior variance decreases because the net gain in flow knowledge by observing near a saddle point outweighs the net loss by the control polluting the observations.

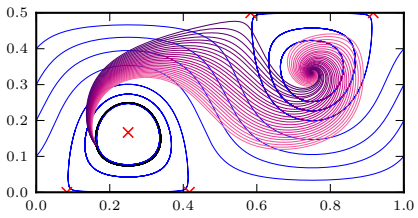

Figure 9: The true drifter paths for each value of ζ for the experiments shown in figure 7(a). The pink path corresponds to the magntidue $\zeta = 0.3$ and the purple path corresponds to $\zeta = 0.55$. The posterior variance decays as ζ approaches 0.55.

5 Conclusion

To summarise, we have measured the performance of two naïve control methods, and one a posteriori control method, both in a time-independent and time-dependent flow. We have done so by observing their influence on the posterior variance in the mean flow direction. Section 3 addresses the naïve controls and section 4 the a posteriori control. Each control is designed to push ocean drifters into uncharted flow regimes. The three cases of control we employ here are a purely zonal control; a control of equal magnitude in both the x and y directions; and the gradient of the posterior mean constructed using a posteriori information from a previous Bayesian update. In the time-independent flow, we show a sizeable reduction of the posterior variance in the mean flow direction for these three cases of control. We also see that on comparing the posterior variance for the zonal and bi-directional controls, similar structures arise when viewed as a function of control magnitude, which dictates when the drifter leaves the eddy and is the main influence on the posterior information. In the case of the a posteriori control in the time-independent flow, the drifter leaves the eddy for all the values of control magnitude we have chosen. Here we observe the variance reduction occurring when the true drifter path approaches a hyperbolic fixed point on the transport barrier of the eddy in the upper-right of the domain. This is evidence that oceanic transport barriers heavily influence posterior information and sets up a novel geometric correspondence between the flow structure and the posterior variance. Using the naïve controls in the time-dependent flow, we show *robustness* of posterior variance as a function of the perturbation parameter. When the control magnitude is such that the drifter leaves the eddy in the *unperturbed* flow, we see reduction in the posterior variance on the initial condition for the time-periodic flow. When employing a time-dependent a posteriori control, we see no overall net gain in posterior variance over the uncontrolled case. For our particular flow and drifter initial condition, it is the case that the uncontrolled drifter path explores a hyperbolic fixed point of an eddy in the time-dependent flow more effectively than the controlled path. This reiterates the efficacy of control strategies and their influence on the path along which observations are made.

There are a number of ways in which this work could be generalised in order to obtain a deeper understanding of the effects controlled ocean drifters have on flow uncertainty. For example, (i) the study of non-periodic model dynamics; (ii) the use of information from the posterior *variance*; (iii) more elaborate control strategies. Many other generalisations are also possible. Non-periodic models are more dynamically consistent with regards to their approximation of larger ocean models. We have seen the application of posterior knowledge in the construction of a control, though only through use of the mean. The variance of the underlying flow could be used in a similar fashion, perhaps to control ocean drifters towards an area of large variance. This could have a similar affect on the posterior distribution as the method of controlling a drifter into a new, unexplored flow regime. Moreover, controls could be constructed to better reflect reality. Ocean gliders have a limited amount of battery power.

Utilising this knowledge in designing a mission plan to optimise a glider's lifespan certainly has its practical applications. Controls that minimise the pollution of the observed data is also desirable. Throughout this paper, we have only used information from one previous Bayesian update. Constructing and executing a posteriori control strategies is a paradigm well suited to that of a Kalman or particle filter; updating the control every time an analysis step is performed. This is left for future discussion.

6 Acknowledgements

Author McDougall would like to acknowledge the work of John Hunter (1968–2012), who led the development of an open-source and freely available plotting library, matplotlib, capable of producing publication-quality graphics [23]. All the figures in this publication were produced with matplotlib.

References

[1] J L Anderson. A Local Least Squares Framework for Ensemble Filtering. *Monthly Weather Review*, 131(4):634–642, April 2003.

[2] A Apte, M Hairer, A M Stuart, and J Voss. Sampling the posterior: An approach to non-Gaussian data assimilation. *Physica D: Nonlinear Phenomena*, 230(1-2):50–64, June 2007.

[3] A Apte, C K R T Jones, and A M Stuart. A Bayesian approach to Lagrangian data assimilation. *Tellus A*, 60(2):336–347, March 2008.

[4] A Apte, C K R T Jones, A M Stuart, and J Voss. Data assimilation: Mathematical and statistical perspectives. *International Journal for Numerical Methods in Fluids*, 56:1033–1046, 2008.

[5] Y F Atchadé. An Adaptive Version for the Metropolis Adjusted Langevin Algorithm with a Truncated Drift. *Methodology and Computing in Applied Probability*, 8(2):235–254, August 2006.

[6] Y F Atchadé and J S Rosenthal. On adaptive Markov chain Monte Carlo algorithms. *Bernoulli*, 11(5):815–828, October 2005.

[7] R W Barbieri and P S Schopf. Oceanographic applications of the Kalman filter. 1982.

[8] L Bengtsson. 4-dimensional assimilation of meteorological observations. *World Meteorological Organization*, 1975.

[9] A Beskos, G O Roberts, and A M Stuart. Optimal scalings for local Metropolis-Hastings chains on nonproduct targets in high dimensions. *The Annals of Applied Probability*, 19(3):863–898, June 2009.

[10] V I Bogachev. *Gaussian Measures*. American Mathematical Society, 1998.

[11] E F Carter. Assimilation of Lagrangian data into a numerical model. *Dynamics of Atmospheres and Oceans*, 13(3-4):335–348, 1989.

[12] S L Cotter, M Dashti, J C Robinson, and A M Stuart. Bayesian inverse problems for functions and applications to fluid mechanics. *Inverse Problems*, 25(11):115008, November 2009.

[13] S L Cotter, M Dashti, and A M Stuart. Approximation of Bayesian inverse problems for PDEs. *SIAM Journal of Numerical Analysis*, 48(1):322–345, 2010.

[14] S L Cotter, M Dashti, and A M Stuart. Variational data assimilation using targetted random walks. *International Journal for Numerical Methods in Fluids*, 68:403–421, 2011.

[15] S L Cotter, G O Roberts, A M Stuart, and D White. MCMC Methods for functions: Modifying old algorithms to make them faster. 2012.

[16] P Courtier, J-N Thépaut, and A Hollingsworth. A strategy for operational implementation of 4D-Var, using an incremental approach. *Quarterly Journal of the Royal Meteorological Society*, 120(519):1367–1387, 1994.

[17] A Doucet, N de Freitas, and N Gordon. *Sequential Monte Carlo Methods In Practice*. 2001.

[18] G Evensen. *Data Assimilation: The Ensemble Kalman Filter*. Springer, 2006.

[19] W K Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

[20] R Herbei and I McKeague. Hybrid Samplers for Ill-Posed Inverse Problems. *Scandinavian Journal of Statistics*, 36(4):839—-853, 2009.

[21] R Herbei, I W McKeague, and K G Speer. Gyres and Jets: Inversion of Tracer Data for Ocean Circulation Structure. *Journal of Physical Oceanography*, 38(6):1180–1202, June 2008.

[22] P L Houtekamer and H L Mitchell. Data Assimilation Using an Ensemble Kalman Filter Technique. *Monthly Weather Review*, 126:796–811, 1998.

[23] J. D. Hunter. Matplotlib: A 2D Graphics Environment. *Computing in Science and Engineering*, 9(3):90–95, 2007.

[24] J Kaipio and E Somersalo. Statistical inverse problems: Discretization, model reduction and inverse crimes. *Journal of Computational and Applied Mathematics*, 198(2):493–504, January 2007.

[25] J P Kaipio, V Kolehmainen, E Somersalo, and M Vauhkonen. Statistical inversion and Monte Carlo sampling methods in electrical impedance tomography. *Inverse problems*, 16(5):1487, 2000.

[26] R E Kalman. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(Series D):35–45, 1960.

[27] R E Kalman and R S Bucy. New results in linear filtering and prediction theory. *Journal of Basic Engineering*, 83:95–107, 1961.

[28] E Kalnay. *Atmospheric modeling, data assimilation and predictability.* Cambridge University Press, 2002.

[29] L Kuznetsov, K Ide, and C K R T Jones. A Method for Assimilation of Lagrangian Data. *Monthly Weather Review*, 131:2247–2260, 2003.

[30] A S Lawless, S Gratton, and N K Nichols. An investigation of incremental 4D-Var using non-tangent linear models. Technical report, 2005.

[31] A S Lawless, S Gratton, and N K Nichols. Approximate iterative methods for variational data assimilation. Technical Report April, 2005.

[32] F-X Le Dimet and O Talagrand. Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects. *Tellus A*, (38A):97–110, 1986.

[33] W Lee, D McDougall, and A M Stuart. Kalman filtering and smoothing for linear wave equations with model error. *Inverse Problems*, 27(9):095008, September 2011.

[34] P J V Leeuwen. Nonlinear data assimilation in geosciences: an extremely efficient particle filter. *Quarterly Journal of the Royal Meteorological Society*, 136(653):1991–1999, October 2010.

[35] J M Lewis and J C Derber. The use of adjoint equations to solve a variational adjustment problem with advective constraints. *Tellus A*, (37A):309–322, 1985.

[36] A C Lorenc. Analysis methods for numerical weather prediction. *Quarterly Journal of the Royal Meteorological Society*, 112:1177–1194, 1986.

[37] A C Lorenc, S P Ballard, R S Bell, N B Ingleby, P L F Andrews, D M Barker, J R Bray, A M Clayton, T Dalby, D Li, T J Payne, and F W Saunders. The Met. Office global three-dimensional variational data assimilation scheme. *Quarterly Journal of the Royal Meteorological Society*, 126(570):2991–3012, October 2000.

[38] I W McKeague, G Nicholls, K Speer, and R Herbei. Statistical inversion of South Atlantic circulation in an abyssal neutral density layer. *Journal of Marine Research*, 63(4):683–704, July 2005.

[39] N Metropolis, A W Rosenbluth, M N Rosenbluth, A H Teller, and E Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*, 21(6):1087, 1953.

[40] A M Michalak. A method for enforcing parameter nonnegativity in Bayesian inverse problems with an application to contaminant source identification. *Water Resources Research*, 39(2):1–14, 2003.

[41] R N Miller. Toward the Application of the Kalman Filter to Regional Open Ocean Modeling. *Journal of Physical Oceanography*, 16:72–86, 1986.

[42] K Mosegaard and A Tarantola. Monte Carlo sampling of solutions to inverse problems. *Journal of Geophysical Research*, 100(B7):12431—-12447, 1995.

[43] D F Parrish and S E Cohn. A Kalman filter for a two-dimensional shallow-water model. In *7th Conference on Numerical Weather Prediction, Montreal, Canada*, pages 1–8, 1985.

[44] R. T. Pierrehumbert. Chaotic Mixing of Tracer and Vorticity by Modulated Traveling Rossby Waves. *Geophysical and Astrophysical Fluid Dynamics*, 58:285–320, 1991.

[45] A A Robel, M Susan Lozier, S F Gary, G L Shillinger, H Bailey, and S J Bograd. Projecting uncertainty onto marine megafauna trajectories. *Deep Sea Research Part I: Oceanographic Research Papers*, 58(12):915–921, September 2011.

[46] G O Roberts. Weak convergence and optimal scaling of random walk Metropolis Algorithms. *Annals of Applied Probability*, 7(1):110–120, 1997.

[47] G O Roberts and J S Rosenthal. Optimal scaling of discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1):255–268, February 1998.

[48] G O Roberts and J S Rosenthal. Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science*, 16(4):351–367, November 2001.

[49] A R Robinson and D B Haidvogel. Dynamical Forceast Experiments with a Baroptropic Open Ocean Model. *Journal of Physical Oceanography*, 10:1928, 1981.

[50] A R Robinson and W G Leslie. Estimation and Prediction of Oceanic Eddy Fields. *Progress in Oceanography*, 14:485–510, 1985.

[51] D L Rudnick, R E Davis, C C Eriksen, D M Fratantoni, and M J Perry. Underwater Gliders for Ocean Research. *Marine Technology Society Journal*, 38(2):73–84, June 2004.

[52] H W Sorenson. *Kalman filtering: theory and application*. IEEE, 1960.

[53] A M Stuart. Inverse Problems: A Bayesian Perspective. In *Acta Numerica*, pages 1–107. 2010.

[54] O Talagrand and P Courtier. Variational assimilation of meteorological observations with the adjoint vorticity equation. I: Theory. *Quarterly Journal of the Royal Meteorological Society*, 113:1311–1328, 1987.